

Concentration of measure without independence: a unified approach via the martingale method

Aryeh Kontorovich* Maxim Raginsky†

August 4, 2016

Abstract

The concentration of measure phenomenon may be summarized as follows: a function of many weakly dependent random variables that is not too sensitive to any of its individual arguments will tend to take values very close to its expectation. This phenomenon is most completely understood when the arguments are mutually independent random variables, and there exist several powerful complementary methods for proving concentration inequalities, such as the martingale method, the entropy method, and the method of transportation inequalities. The setting of dependent arguments is much less well understood. This chapter focuses on the martingale method for deriving concentration inequalities without independence assumptions. In particular, we use the machinery of so-called Wasserstein matrices to show that the Azuma-Hoeffding concentration inequality for martingales with almost surely bounded differences, when applied in a sufficiently abstract setting, is powerful enough to recover and sharpen several known concentration results for nonproduct measures. Wasserstein matrices provide a natural formalism for capturing the interplay between the metric and the probabilistic structures, which is fundamental to the concentration phenomenon.

1 Introduction

At its most abstract, the concentration of measure phenomenon may be summarized as follows: a function of several weakly dependent random variables that is not too sensitive to any of the individual arguments will tend to take values very close to its expectation. This phenomenon is most completely understood in the case of independent arguments, and the recent book [2] provides an excellent survey (see also [30] for an exposition from the viewpoint of, and with applications to, information theory).

The case of dependent arguments has yet to mature into such a unified, overarching theory. The earliest concentration results for non-product measures were established for Haar measures on various groups, and relied strongly on the highly symmetric nature of the Haar measure in question. These results include Lévy's classic isoperimetric inequality on the sphere [20] and Maurey's concentration inequality on the permutation group [28]. To the best of our knowledge, the first concentration result for a non-product, non-Haar measure is due to Marton [22], where she proved a McDiarmid-type bound for contracting Markov chains. A flurry of activity followed. Besides

*Department of Computer Science, Ben-Gurion University, Beer-Sheva, Israel. E-mail: karyeh@cs.bgu.ac.il.

†Department of Electrical and Computer Engineering and the Coordinated Science Laboratory, University of Illinois, Urbana, IL, USA. E-mail: maxim@illinois.edu.

Marton's own follow-up work [23, 24, 25], the transportation method she pioneered was extended by Samson [33], and martingale techniques [32, 6, 17], as well as methods relying on the Dobrushin interdependence matrix [19, 4, 36], have been employed in obtaining concentration results for non-product measures. The underlying theme is that the independence assumption may be relaxed to one of *weak* dependence, the latter being quantified by various mixing coefficients.

This chapter is an attempt at providing an abstract unifying framework that generalizes and sharpens some of the above results. This framework combines classical martingale techniques with the method of *Wasserstein matrices* [10]. In particular, we rely on Wasserstein matrices to obtain general-purpose quantitative estimates of the local variability of a function of many dependent random variables after taking a conditional expectation with respect to a subset of the variables. A concentration inequality in a metric space must necessarily capture the interplay between the metric and the distribution, and, in our setting, Wasserstein matrices provide the ideal analytical tool for this task. As an illustration, we recover (and, in certain cases, sharpen) some results of [19, 6, 17] by demonstrating all of these to be special cases of the Wasserstein matrix method.

The remainder of the chapter is organized as follows. Section 2 is devoted to setting up the basic notation and preliminary definitions. A brief discussion of the concentration of measure phenomenon in high-dimensional spaces is presented in Section 3, together with a summary of key methods to establish concentration under the independence assumption. Next, in Section 4, we present our abstract martingale technique and then demonstrate its wide scope in Section 5 by deriving many of previously published concentration inequalities as special cases. We conclude in Section 6 by listing some open questions.

2 Preliminaries and notation

2.1 Metric probability spaces

A *metric probability space* is a triple (Ω, μ, d) , where Ω is a Polish space equipped with its Borel σ -field, μ is a Borel probability measure on Ω , and d is metric on Ω which is a measurable function on the product space $\Omega \times \Omega$. We do not assume that d is the same metric that metrizes the Polish topology on Ω .

2.2 Product spaces

Since concentration of measure is a high-dimensional phenomenon, a natural setting for studying it is that of a product space. Let T be a finite index set, which we identify with the set $[n] \triangleq \{1, \dots, n\}$, where $n = |T|$ (this amounts to fixing some linear ordering of the elements of T). We will use the following notation for subintervals of T : $[i] \triangleq \{1, \dots, i\}$; $[i, j] \triangleq \{i, i + 1, \dots, j\}$ for $i \neq j$; $(i, j] \triangleq \{i + 1, \dots, j\}$ for $i < j$; $(i, j) \triangleq \{i + 1, \dots, j - 1\}$ for $i < j - 1$; etc.

With each $i \in T$, we associate a measurable space $(\mathsf{X}_i, \mathcal{B}_i)$, where X_i is a Polish space and \mathcal{B}_i is its Borel σ -field. For each $I \subseteq T$, we will equip the product space $\mathsf{X}^I \triangleq \prod_{i \in I} \mathsf{X}_i$ with the product σ -field $\mathcal{B}^I \triangleq \bigotimes_{i \in I} \mathcal{B}_i$. When $I = T$, we will simply write X and \mathcal{B} . We will write x^I and x for a generic element of X^I and X , respectively. Given two sets $I, J \subseteq T$ with $I \cap J = \emptyset$, the

concatenation of $x^I \in \mathsf{X}^I$ and $z^J \in \mathsf{X}^J$ is defined as $y = x^I z^J \in \mathsf{X}^{I \cup J}$ by setting

$$y_i = \begin{cases} x_i, & i \in I \\ z_i, & i \in J \end{cases}.$$

Given a random object $X = (X_i)_{i \in T}$ taking values in X according to a probability law μ , we will denote by $\mathbb{P}_\mu[\cdot]$ and $\mathbb{E}_\mu[\cdot]$ the probability and expectation with respect to μ , by $\mu^I(dx^I|x^J)$ the regular conditional probability law of X^I given $X^J = x^J$, and by $\mu^I(dx^I)$ the marginal probability law of X^I . When $I = \{i\}$, we will write $\mu_i(\cdot)$ and $\mu_i(\cdot|x^J)$.

For each $i \in T$, we fix a metric on X_i , which is assumed to be measurable with respect to the product σ -field $\mathcal{B}_i \otimes \mathcal{B}_i$. For each $I \subseteq T$, equip X^I with the product metric ρ^I , where

$$\rho^I(x^I, z^I) \triangleq \sum_{i \in I} \rho_i(x_i, z_i), \quad \forall x^I, z^I \in \mathsf{X}^I.$$

When $I \equiv T$, we will simply write ρ instead of ρ^T . In this way, for any Borel probability measure μ on X , we can introduce a ‘‘global’’ metric probability space (X, μ, ρ) , as well as ‘‘local’’ metric probability spaces $(\mathsf{X}^I, \mu^I, \rho^I)$ and $(\mathsf{X}^I, \mu^I(\cdot|x^J), \rho_I)$ for all $I, J \subseteq T$ and all $x^J \in \mathsf{X}^J$.

2.3 Couplings and transportation distances

Let Ω be a Polish space. A *coupling* of two Borel probability measures μ and ν on Ω is a Borel probability measure \mathbf{P} on the product space $\Omega \times \Omega$, such that $\mathbf{P}(\cdot \times \Omega) = \mu$ and $\mathbf{P}(\Omega \times \cdot) = \nu$. We denote the set of all couplings of μ and ν by $\mathcal{C}(\mu, \nu)$. Let d be a lower-semicontinuous metric on Ω . We denote by $\text{Lip}(\Omega, d)$ the space of all functions $\Omega \rightarrow \mathbb{R}$ that are Lipschitz with respect to d , and by $\text{Lip}_c(\Omega, d)$ the subset of $\text{Lip}(\Omega, d)$ consisting of c -Lipschitz functions. The L^1 *Wasserstein* (or *transportation*) *distance* between μ and ν is defined as

$$W_d(\mu, \nu) \triangleq \inf_{\mathbf{P} \in \mathcal{C}(\mu, \nu)} \mathbb{E}_{\mathbf{P}}[d(X, Y)], \quad (2.1)$$

where (X, Y) is a random element of $\Omega \times \Omega$ with law \mathbf{P} . The transportation distance admits a dual (Kantorovich–Rubinstein) representation

$$W_d(\mu, \nu) = \sup_{f \in \text{Lip}_1(\Omega, d)} \left| \int_{\Omega} f d\mu - \int_{\Omega} f d\nu \right|. \quad (2.2)$$

For example, when we equip Ω with the trivial metric $d(\omega, \omega') = \mathbf{1}\{\omega \neq \omega'\}$, the corresponding Wasserstein distance coincides with the total variation distance:

$$W_d(\mu, \nu) = \|\mu - \nu\|_{\text{TV}} = \sup_A |\mu(A) - \nu(A)|,$$

where the supremum is over all Borel subsets of Ω .

In the context of the product space (X, ρ) defined earlier, we will use the shorthand W_i for W_{ρ_i} , W^I for W_{ρ^I} , and W for W_{ρ} .

2.4 Markov kernels and Wasserstein matrices

A *Markov kernel* on X is a mapping $K : \mathsf{X} \times \mathcal{B} \rightarrow [0, 1]$, such that $x \mapsto K(x, A)$ is measurable for each $A \in \mathcal{B}$, and $K(x, \cdot)$ is a Borel probability measure on X for each $x \in \mathsf{X}$. Given a Markov kernel K and a bounded measurable function $f : \mathsf{X} \rightarrow \mathbb{R}$, we denote by Kf the bounded measurable function

$$Kf(x) \triangleq \int_{\mathsf{X}} f(y)K(x, dy), \quad x \in \mathsf{X}.$$

Likewise, given a Borel probability measure μ on X , we denote by μK the Borel probability measure

$$\mu K(A) \triangleq \int_{\mathsf{X}} K(x, A)\mu(dx), \quad A \in \mathcal{B}.$$

It is not hard to see that $\int_{\mathsf{X}} f d(\mu K) = \int_{\mathsf{X}} (Kf) d\mu$.

Given a measurable function $f : \mathsf{X} \rightarrow \mathbb{R}$, we define the *local oscillation of f at $i \in T$* as

$$\delta_i(f) \triangleq \sup_{\substack{x, z \in \mathsf{X} \\ x^{T \setminus \{i\}} = z^{T \setminus \{i\}}}} \frac{|f(x) - f(z)|}{\rho_i(x_i, z_i)},$$

where we follow the convention $0/0 = 0$. This quantity measures the variability of f in its i th argument when all other arguments are held fixed. As will become evident later on, our martingale technique for establishing concentration inequalities for a given function $f : \mathsf{X} \rightarrow \mathbb{R}$ requires controlling the local oscillations $\delta_i(Kf)$ in terms of the local oscillations $\delta_i(f)$ for appropriately chosen Markov kernels K .

To get an idea of what is involved, let us consider the simple case when each X_i is endowed with the scaled trivial metric $\rho_i(x_i, z_i) \triangleq \alpha_i \mathbf{1}\{x_i \neq z_i\}$, where $\alpha_i > 0$ is some fixed constant. Then

$$\delta_i(f) = \frac{1}{\alpha_i} \sup \left\{ |f(x) - f(z)| : x, z \in \mathsf{X}, x^{T \setminus \{i\}} = z^{T \setminus \{i\}} \right\}.$$

The corresponding metric ρ on X is the weighted Hamming metric

$$\rho_{\alpha}(x, z) \triangleq \sum_{i \in T} \alpha_i \mathbf{1}\{x_i \neq z_i\}. \quad (2.3)$$

Fix a Markov kernel K on X . The *Dobrushin contraction coefficient* of K (also associated in the literature with Doeblin's name) is the smallest $\theta \geq 0$ for which $\|K(x, \cdot) - K(z, \cdot)\|_{\text{TV}} \leq \theta$ holds for all $x, z \in \mathsf{X}$. The term *contraction* is justified by the well-known inequality (apparently going back to Markov himself [21, §5])

$$\|\mu K - \nu K\|_{\text{TV}} \leq \theta \|\mu - \nu\|_{\text{TV}}, \quad (2.4)$$

which holds for all probability measures μ, ν on X . Then we have the following estimate:

Proposition 2.1 *If K is a Markov kernel on X with Dobrushin coefficient θ , then for every $i \in T$ and for every $f \in \text{Lip}(\mathsf{X}, \rho_{\alpha})$, we have*

$$\delta_i(Kf) \leq \frac{\theta}{\alpha_i} \sum_{j \in T} \alpha_j \delta_j(f). \quad (2.5)$$

Proof: Fix an index $i \in T$ and any two $x, z \in \mathsf{X}$ that differ only in the i th coordinate: $x^{T \setminus \{i\}} = z^{T \setminus \{i\}}$ and $x_i \neq z_i$. Pick an arbitrary coupling $\mathbf{P}_{x,z} \in \mathcal{C}(K(x, \cdot), K(z, \cdot))$. Then

$$\begin{aligned}
|Kf(x) - Kf(z)| &= \left| \int_{\mathsf{X}} K(x, du) f(u) - \int_{\mathsf{X}} K(z, dy) f(y) \right| \\
&= \left| \int_{\mathsf{X} \times \mathsf{X}} \mathbf{P}_{x,z}(du, dy) (f(u) - f(y)) \right| \\
&\leq \sum_{j \in T} \delta_j(f) \int_{\mathsf{X} \times \mathsf{X}} \mathbf{P}_{x,z}(du, dy) \rho_j(u_j, y_j) \\
&= \sum_{j \in T} \alpha_j \delta_j(f) \int_{\mathsf{X} \times \mathsf{X}} \mathbf{P}_{x,z}(du, dy) \mathbf{1}\{u_j \neq y_j\} \\
&\leq \sum_{j \in T} \alpha_j \delta_j(f) \cdot \int_{\mathsf{X} \times \mathsf{X}} \mathbf{P}_{x,z}(du, dy) \mathbf{1}\{u \neq y\},
\end{aligned}$$

where the first inequality is by the definition of $\delta_i(f)$, while the second one follows from the obvious implication $u_j \neq y_j \Rightarrow u \neq y$. Taking the infimum of both sides over all couplings $\mathbf{P}_{x,z} \in \mathcal{C}(K(x, \cdot), K(z, \cdot))$ yields

$$\begin{aligned}
|Kf(x) - Kf(z)| &\leq \sum_{j \in T} \alpha_j \delta_j(f) \cdot \|K(x, \cdot) - K(z, \cdot)\|_{\text{TV}} \\
&\leq \theta \sum_{j \in T} \alpha_j \delta_j(f).
\end{aligned}$$

Finally, dividing both sides of the above inequality by α_i and taking the supremum over all choices of x, z that differ only in the i th coordinate, we obtain (2.5). \blacksquare

One shortcoming of the above result (which is nontrivial only under the rather strong condition

$$\theta < \frac{\alpha_i}{\alpha_j} < \theta^{-1} \tag{2.6}$$

for all $i, j \in T$) is that it gives only a very rough idea of the influence of $\delta_j(f)$ for $j \in T$ on $\delta_i(Kf)$. For example, if $\alpha_1 = \dots = \alpha_n = 1$, then the condition (2.6) reduces to the Dobrushin contraction condition $\theta < 1$, and the inequality (2.5) becomes

$$\delta_i(Kf) \leq \theta \sum_{j \in T} \delta_j(f),$$

suggesting that all of the $\delta_j(f)$'s influence $\delta_i(Kf)$ equally. However, this picture can be refined. To that end, we introduce the notion of a *Wasserstein matrix* following Föllmer [10]. Let us denote by $\boldsymbol{\delta}(f)$ the vector $(\delta_i(f))_{i \in T}$. We say that a nonnegative matrix $V = (V_{ij})_{i,j \in T}$ is a Wasserstein matrix for K if, for every $f \in \text{Lip}(\mathsf{X}, \rho)$ and for every $i \in T$,

$$\delta_i(Kf) \leq \sum_{j \in T} V_{ij} \delta_j(f), \tag{2.7}$$

or, in vector form, if $\boldsymbol{\delta}(Kf) \preceq V \boldsymbol{\delta}(f)$.

One of our main objectives will be to show that concentration inequalities for functions f of $X \sim \mu$ can be obtained using Wasserstein matrices for certain Markov kernels K related to μ . In order to motivate the introduction of Wasserstein matrices, we record a couple of contraction estimates for Markov kernels that may be of independent interest. To that end, we introduce another coupling-based distance between probability measures [2, Chap. 8]: for two Borel probability measures on X , define

$$\bar{W}(\mu, \nu) \triangleq \inf_{\mathbf{P} \in \mathcal{C}(\mu, \nu)} \sqrt{\sum_{i \in T} (\mathbb{E}_{\mathbf{P}}[\rho_i(X_i, Y_i)])^2}, \quad (2.8)$$

where (X, Y) is a random element of $\mathsf{X} \times \mathsf{X}$. Even though \bar{W} is not a Wasserstein distance, we can use the inequality $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for $a, b \geq 0$ to show that $\bar{W}(\mu, \nu) \leq W(\mu, \nu)$.

Proposition 2.2 *Let V be a Wasserstein matrix for a Markov kernel K on X . Then for any Lipschitz function $f : \mathsf{X} \rightarrow \mathbb{R}$,*

$$|\mathbb{E}_{\mu K}[f(X)] - \mathbb{E}_{\nu K}[f(X)]| \leq \|V\boldsymbol{\delta}(f)\|_{\ell^2(T)} \bar{W}(\mu, \nu). \quad (2.9)$$

Proof: Fix an arbitrary coupling $\mathbf{P} \in \mathcal{C}(\mu, \nu)$ and let (X, Y) be a random element of $\mathsf{X} \times \mathsf{X}$ with law \mathbf{P} . Then

$$\begin{aligned} |\mathbb{E}_{\mu K}[f(X)] - \mathbb{E}_{\nu K}[f(X)]| &= |\mathbb{E}_{\mu}[Kf(X)] - \mathbb{E}_{\nu}[Kf(X)]| \\ &= |\mathbb{E}_{\mathbf{P}}[Kf(X) - Kf(Y)]| \\ &\leq \sum_{i \in T} \delta_i(Kf) \cdot \mathbb{E}_{\mathbf{P}}[\rho_i(X_i, Y_i)] \\ &\leq \sum_{i \in T} \sum_{j \in T} V_{ij} \delta_j(f) \cdot \mathbb{E}_{\mathbf{P}}[\rho_i(X_i, Y_i)]. \end{aligned}$$

where in the last step we have used the definition of the Wasserstein matrix. Using the Cauchy–Schwarz inequality, we obtain

$$\begin{aligned} |\mathbb{E}_{\mu K}[f(X)] - \mathbb{E}_{\nu K}[f(X)]| &\leq \sqrt{\sum_{i \in T} \left| \sum_{j \in T} V_{ij} \delta_j(f) \right|^2 \cdot \sum_{i \in T} (\mathbb{E}_{\mathbf{P}}[\rho_i(X_i, Y_i)])^2} \\ &= \|V\boldsymbol{\delta}(f)\|_{\ell^2(T)} \cdot \sqrt{\sum_{i \in T} (\mathbb{E}_{\mathbf{P}}[\rho_i(X_i, Y_i)])^2}. \end{aligned}$$

Taking the infimum of both sides over all $\mathbf{P} \in \mathcal{C}(\mu, \nu)$, we obtain (2.9). ■

Corollary 2.3 *Let V be a Wasserstein matrix for a Markov kernel K on X . Then, for any two Borel probability measures μ and ν on X ,*

$$W(\mu K, \nu K) \leq \|V\mathbf{1}\|_{\ell^2(T)} \bar{W}(\mu, \nu), \quad (2.10)$$

where $\mathbf{1} \in \mathbb{R}^T$ is the vector of all ones, and therefore

$$W(\mu K, \nu K) \leq \|V\mathbf{1}\|_{\ell^2(T)} W(\mu, \nu).$$

Proof: A function $f : \mathsf{X} \rightarrow \mathbb{R}$ belongs to $\text{Lip}_1(\mathsf{X}, \rho)$ if and only if $\delta(f) \in [0, 1]^T$. Using the dual representation (2.2) of W and applying Proposition 2.2, we can write

$$\begin{aligned} W(\mu K, \nu K) &= \sup_{f \in \text{Lip}_1(\mathsf{X}, \rho)} |\mathbb{E}_{\mu K}[f(X)] - \mathbb{E}_{\nu K}[f(X)]| \\ &\leq \sup_{\xi \in [0, 1]^T} \|V\xi\|_{\ell^2(T)} \bar{W}(\mu, \nu). \end{aligned}$$

Since V is a nonnegative matrix, the supremum is achieved by $\xi = \mathbf{1}$. ■

2.5 Relative entropy

Finally, we will need some key notions from information theory. The *relative entropy* (or *information divergence*) between two probability measures μ, ν on a space Ω is defined as

$$D(\nu \parallel \mu) \triangleq \begin{cases} \int_{\Omega} d\mu f \log f, & \text{if } \nu \ll \mu \text{ with } f = d\nu/d\mu \\ +\infty, & \text{otherwise} \end{cases}.$$

We use natural logarithms throughout the chapter. The relative entropy is related to the total variation distance via Pinsker's inequality¹

$$\|\mu - \nu\|_{\text{TV}} \leq \sqrt{\frac{1}{2} D(\mu \parallel \nu)}. \quad (2.11)$$

3 Concentration of measure and sufficient conditions

In this section, we give a precise definition of the concentration of measure phenomenon, review several sufficient conditions for it to hold, and briefly discuss how it can be established under the independence assumption via tensorization. For more details and further references, the reader can consult [2] or [30].

We say that the metric probability space (X, μ, ρ) has the concentration of measure property if there exists a positive constant $c > 0$, such that, for every Lipschitz function $f : \mathsf{X} \rightarrow \mathbb{R}$,

$$\mathbb{P}_{\mu} \{f(X) - \mathbb{E}_{\mu}[f(X)] \geq t\} \leq e^{-t^2/2c\|f\|_{\text{Lip}}^2}, \quad \forall t > 0 \quad (3.12)$$

where

$$\|f\|_{\text{Lip}} \triangleq \sup_{\substack{x, y \in \mathsf{X} \\ x \neq y}} \frac{|f(x) - f(y)|}{\rho(x, y)}$$

is the Lipschitz constant of f . A sufficient (and, up to constants, necessary) condition for (3.12) is that, for every $f \in \text{Lip}_1(\mathsf{X}, \rho)$, the random variable $f(X)$ with $X \sim \mu$ is *c-subgaussian*, i.e.,

$$\log \mathbb{E}_{\mu} \left[e^{\lambda(f(X) - \mathbb{E}_{\mu}[f(X)])} \right] \leq \frac{c\lambda^2}{2}, \quad \forall \lambda \in \mathbb{R}. \quad (3.13)$$

¹ Though commonly referred to as *Pinsker's inequality*, (2.11) as given here (with the optimal constant $\frac{1}{2}$) was proven by Csiszár [7] and Kullback [18] in 1967.

A fundamental result of Bobkov and Götze [1] states that the subgaussian estimate (3.13) holds for all $f \in \text{Lip}_1(\mathsf{X}, \rho)$ if and only if μ satisfies the so-called *transportation-information inequality*

$$W(\mu, \nu) \leq \sqrt{2cD(\nu\|\mu)}, \quad (3.14)$$

where ν ranges over all Borel probability measures on X . We will use the shorthand $\mu \in T_\rho(c)$ to denote the fact that the inequality (3.14) holds for all ν . The key role of transportation-information inequalities in characterizing the concentration of measure phenomenon was first recognized by Marton in a breakthrough paper [22], with further developments in [23, 24, 25].

The *entropy method* (see, e.g., [2, Chap. 6] and [30, Chap. 3]) provides another route to establishing (3.13). Its underlying idea can be briefly described as follows. Given a measurable function $f : \mathsf{X} \rightarrow \mathbb{R}$, consider the logarithmic moment-generating function

$$\psi_f(\lambda) \triangleq \log \mathbb{E}_\mu \left[e^{\lambda(f(X) - \mathbb{E}_\mu[f(X)])} \right]$$

of the centered random variable $f(X) - \mathbb{E}_\mu[f(X)]$. For any $\lambda \neq 0$, introduce the *tilted probability measure* $\mu^{(\lambda f)}$ with

$$\frac{d\mu^{(\lambda f)}}{d\mu} = \frac{e^{\lambda f}}{\mathbb{E}_\mu[e^{\lambda f}]} = \frac{e^{\lambda(f - \mathbb{E}_\mu f)}}{e^{\psi_f(\lambda)}}.$$

Then a simple calculation shows that the relative entropy $D(\mu^{(\lambda f)}\|\mu)$ can be expressed as

$$D(\mu^{(\lambda f)}\|\mu) = \lambda\psi_f'(\lambda) - \psi_f(\lambda) \equiv \lambda^2 \left(\frac{\psi_f(\lambda)}{\lambda} \right)'$$

where the prime denotes differentiation with respect to λ . Using the fact that $\psi_f(0) = 0$ and integrating, we obtain the following formula for $\psi_f(\lambda)$:

$$\psi_f(\lambda) = \lambda \int_0^\lambda \frac{D(\mu^{(tf)}\|\mu)}{t^2} dt. \quad (3.15)$$

This representation is at the basis of the so-called *Herbst argument*, which for our purposes can be summarized as follows:

Lemma 3.1 (Herbst) *The metric probability space (X, μ, ρ) has the concentration property with constant c if, for any $f \in \text{Lip}_1(\mathsf{X}, \rho)$,*

$$D(\mu^{(tf)}\|\mu) \leq \frac{ct^2}{2}, \quad \forall t > 0. \quad (3.16)$$

Remark 3.1 Up to a constant, the converse is also true [34, Prob. 3.12]: if the subgaussian estimate (3.13) holds for every $f \in \text{Lip}_1(\mathsf{X}, \rho)$, then

$$D(\mu^{(tf)}\|\mu) \leq 2ct^2, \quad \forall t > 0$$

for every $f \in \text{Lip}_1(\mathsf{X}, \rho)$.

In this way, the problem of establishing the concentration phenomenon reduces to showing that (3.16) holds for every $f \in \text{Lip}_1(\mathsf{X}, \rho)$, typically via logarithmic Sobolev inequalities or other functional inequalities.

3.1 Concentration of measure under the independence assumption

To set the stage for the general treatment of the concentration phenomenon in high dimensions, we first consider the independent case, i.e., when coordinates X_i , $i \in T$, of the random object $X \sim \mu$ are mutually independent. In other words, the probability measure μ is equal to the product of its marginals: $\mu = \mu_1 \otimes \dots \otimes \mu_n$. The key to establishing the concentration property in such a setting is *tensorization*, which is an umbrella term for any result that allows one to derive the “global” concentration property of the high-dimensional product space $(\mathsf{X}_1 \otimes \dots \otimes \mathsf{X}_n, \rho, \mu_1 \otimes \dots \otimes \mu_n)$ from “local” concentration properties of the coordinate spaces $(\mathsf{X}_i, \rho_i, \mu_i)$, $i \in T$.

Below, we list two such tensorization results, one for the transportation-information inequalities and one for the relative entropy. Both of these results are deep consequences of the interplay between the independence structure of μ and the metric structure of ρ . Indeed, a function $f : \mathsf{X} \rightarrow \mathbb{R}$ belongs to $\text{Lip}_1(\mathsf{X}, \rho)$ if and only if $\delta_i(f) \leq 1$ for all $i \in T$, i.e., if and only if, for every $i \in T$ and every $x^{T \setminus \{i\}} \in \mathsf{X}^{T \setminus \{i\}}$, the function $f_i : \mathsf{X}_i \rightarrow \mathbb{R}$ given by $f_i(y_i) \triangleq f(y_i x^{T \setminus \{i\}})$ is 1-Lipschitz with respect to the metric ρ_i on X_i . With this in mind, it is reasonable to expect that if one can establish a concentration property for all 1-Lipschitz functions on the coordinate spaces X_i , then one can deduce a concentration property for functions on the product space X that are 1-Lipschitz in each coordinate.

Lemma 3.2 (Tensorization of transportation-information inequalities) *Suppose that there exist constants $c_1, \dots, c_n \geq 0$, such that*

$$\mu_i \in T_{\rho_i}(c_i), \quad \forall i \in T.$$

Then $\mu = \mu_1 \otimes \dots \otimes \mu_n \in T_\rho(c)$ with $c = \sum_{i=1}^n c_i$.

For example, by an appropriate rescaling of Pinsker’s inequality (2.11), we see that, if each coordinate space X_i is endowed with the scaled trivial metric $\rho_i(x_i, z_i) = \alpha_i \mathbf{1}\{x_i \neq z_i\}$ for some $\alpha_i > 0$, then any Borel probability measure μ_i on X_i satisfies the transportation-information inequality with $c_i = \alpha_i^2/4$. By the above tensorization lemma, any product measure $\mu_1 \otimes \dots \otimes \mu_n$ on the product space $\mathsf{X}_1 \otimes \dots \otimes \mathsf{X}_n$ equipped with the weighted Hamming metric ρ_α defined in (2.3) satisfies $T_{\rho_\alpha}(c)$ with $c = \frac{1}{4} \sum_{i \in T} \alpha_i^2$. Consequently, by the Bobkov–Götze theorem, the subgaussian estimate (3.13) holds for any function $f \in \text{Lip}_1(\mathsf{X}, \rho_\alpha)$, which in turn implies, via the Chernoff bound, that

$$\mathbb{P}_\mu \left\{ f - \mathbb{E}_\mu f \geq t \right\} \leq \exp \left(- \frac{2t^2}{\sum_{i \in T} \alpha_i^2} \right), \quad \forall t \geq 0.$$

This provides an alternative derivation of McDiarmid’s inequality (with the sharp constant in the exponent), which was originally proved using the martingale method.

Lemma 3.3 (Tensorization of relative entropy) *Consider a product measure $\mu = \mu_1 \otimes \dots \otimes \mu_n$. Then for any other probability measure ν on X we have*

$$D(\nu \parallel \mu) \leq \sum_{i \in T} \mathbb{E}_\nu D(\nu_i(\cdot | X^{T \setminus \{i\}}) \parallel \mu_i). \quad (3.17)$$

The idea is to apply this lemma to $\nu = \mu^{(tf)}$ for some $t \geq 0$ and an arbitrary $f \in \text{Lip}_1(\mathsf{X}, \rho)$. In that case, a simple calculation shows that the conditional probability measure $\nu_i(dx_i | x^{T \setminus \{i\}}) =$

$\mu_i^{(tf)}(dx_i|x^{T\setminus\{i\}})$ is equal to the tilted distribution $\mu_i^{(tf_i)}$ with $f_i(x_i) = f(x_i x^{T\setminus\{i\}})$, and therefore

$$D(\mu^{(tf)}\|\mu) \leq \sum_{i=1}^n \mathbb{E}_{\mu^{(tf)}} D(\mu_i^{(tf_i)}\|\mu_i).$$

If $f \in \text{Lip}_1(\mathbf{X}, \rho)$, then $f_i \in \text{Lip}_1(\mathbf{X}_i, \rho_i)$. Thus, if we can show that, for any $g \in \text{Lip}_1(\mathbf{X}_i, \rho_i)$,

$$D(\mu_i^{(tg)}\|\mu_i) \leq \frac{c_i t^2}{2}, \quad \forall t \geq 0,$$

then the estimate

$$D(\mu^{(tf)}\|\mu) \leq \frac{ct^2}{2}, \quad \forall t \geq 0$$

holds with $c = \sum_{i=1}^n c_i$ for all $f \in \text{Lip}(\mathbf{X}, \rho)$ by the tensorization lemma. Invoking the Herbst argument, we conclude that (\mathbf{X}, μ, ρ) has the concentration property with the same c .

4 The abstract martingale method

In this section, we present a general martingale-based scheme for deriving concentration inequalities for functions of many dependent random variables. Let $f : \mathbf{X} \rightarrow \mathbb{R}$ be the function of interest, and let $X = (X_i)_{i \in T}$ be a random element of the product space \mathbf{X} with probability law μ . Let $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \dots \subset \mathcal{F}_m$ be a filtration (i.e., an increasing sequence of σ -fields) on \mathbf{X} , such that \mathcal{F}_0 is trivial and $\mathcal{F}_m = \mathcal{B}$. The idea is to decompose the centered random variable $f(X) - \mathbb{E}_\mu[f(X)]$ as a sum of martingale differences

$$M^{(j)} \triangleq \mathbb{E}_\mu[f(X)|\mathcal{F}_j] - \mathbb{E}_\mu[f(X)|\mathcal{F}_{j-1}], \quad j = 1, \dots, m.$$

By construction, $\mathbb{E}_\mu[f(X)|\mathcal{F}_m] = f(X)$ and $\mathbb{E}_\mu[f(X)|\mathcal{F}_0] = \mathbb{E}_\mu[f(X)]$, so the problem of bounding the probability $\mathbb{P}_\mu\{|f - \mathbb{E}_\mu f| \geq t\}$ for a given $t \geq 0$ reduces to bounding the probability

$$\mathbb{P}_\mu\left\{\left|\sum_{j=1}^m M^{(j)}\right| \geq t\right\}.$$

The latter problem hinges on being able to control the martingale differences $M^{(j)}$. In particular, if each $M^{(j)}$ is a.s. bounded, we have the following:

Theorem 4.1 (Azuma–Hoeffding inequality) *Let $\{M^{(j)}\}_{j=1}^m$ be a martingale difference sequence with respect to a filtration $\{\mathcal{F}_j\}_{j=0}^m$. Suppose that, for each j , there exist \mathcal{F}_{j-1} -measurable random variables $A^{(j)}$ and $B^{(j)}$, such that $A^{(j)} \leq M^{(j)} \leq B^{(j)}$ a.s. Then*

$$\mathbb{P}\left[\exp\left(\lambda \sum_{j=1}^m M^{(j)}\right)\right] \leq \exp\left(\frac{\lambda^2 \sum_{j=1}^m \|B^{(j)} - A^{(j)}\|_\infty^2}{8}\right), \quad \forall \lambda \in \mathbb{R}. \quad (4.18)$$

Consequently, for any $t \geq 0$,

$$\mathbb{P}\left\{\left|\sum_{j=1}^m M^{(j)}\right| \geq t\right\} \leq 2 \exp\left(-\frac{2t^2}{\sum_{j=1}^m \|B^{(j)} - A^{(j)}\|_\infty^2}\right). \quad (4.19)$$

The most straightforward choice of the filtration is also the most natural one: take $m = |T| = n$, and for each $i \in T$ take $\mathcal{F}_i = \sigma(X^{[i]})$. For $i \in T$, define a Markov kernel $K^{(i)}$ on \mathbf{X} by

$$K^{(i)}(x, dy) \triangleq \delta_{x^{[i-1]}}(dy^{[i-1]}) \otimes \mu^{[i,n]}(dy^{[i,n]}|x^{[i-1]}). \quad (4.20)$$

Then, for any $f \in L^1(\mu)$ we have

$$\begin{aligned} K^{(i)}f(x) &= \int_{\mathbf{X}} f(y)K^{(i)}(x, dy) \\ &= \int_{\mathbf{X}^{[i,n]}} f(x^{[i-1]}y^{[i,n]})\mu^{[i,n]}(dy^{[i,n]}|x^{[i-1]}) \\ &= \mathbb{E}_\mu[f(X)|X^{[i-1]} = x^{[i-1]}]; \end{aligned}$$

in particular, $K^{(1)}f = \mathbb{E}_\mu f$. We extend this definition to $i = n + 1$ in the obvious way:

$$K^{(n+1)}(x, dy) = \delta_x(dy),$$

so that $K^{(n+1)}f = f$. Then, for each $i \in T$, we can write $M^{(i)} = K^{(i+1)}f - K^{(i)}f$. With this construction, we can state the following theorem that applies to the case when each coordinate space \mathbf{X}_i is endowed with a bounded measurable metric ρ_i :

Theorem 4.2 *Assume that, for all i ,*

$$\|\rho_i\| \triangleq \sup_{x_i, z_i \in \mathbf{X}_i} \rho_i(x_i, z_i) < \infty.$$

For each $i \in \{1, \dots, n+1\}$, let $V^{(i)}$ be a Wasserstein matrix for the Markov kernel $K^{(i)}$ defined in (4.20), in the sense that $\delta(K^{(i)}f) \preceq V^{(i)}\delta(f)$ holds for each $f \in \text{Lip}(\mathbf{X}, \rho)$ as in (2.7). Define the matrix $\Gamma = (\Gamma_{ij})_{i,j \in T}$ with entries

$$\Gamma_{ij} \triangleq \|\rho_i\|V_{ij}^{(i+1)}.$$

Then, for any $f \in \text{Lip}(\mathbf{X}, \rho)$ and for any $t \geq 0$, we have

$$\mathbb{P}_\mu \left\{ |f(X) - \mathbb{E}_\mu[f(X)]| \geq t \right\} \leq 2 \exp \left(- \frac{2t^2}{\|\Gamma\delta(f)\|_{\ell^2(T)}^2} \right). \quad (4.21)$$

Proof: For each $i \in T$, using the tower property of conditional expectations, we can write

$$\begin{aligned} M^{(i)} &= \mathbb{E}_\mu[f(X)|X^{[i]} = x^{[i]}] - \mathbb{E}_\mu[f(X)|X^{[i-1]} = x^{[i-1]}] \\ &= \mathbb{E}_\mu[f(X)|X^{[i]} = x^{[i]}] - \mathbb{E}_\mu[\mathbb{E}_\mu[f(X)|X^{[i-1]} = x^{[i-1]}, X_i]|X^{[i-1]} = x^{[i-1]}] \\ &= \int_{\mathbf{X}^{[i,n]}} \mu^{[i,n]}(dy^{[i,n]}|x^{[i-1]}) \left(\int_{\mathbf{X}^{(i,n)}} \mu^{(i,n)}(dz^{(i,n)}|x^{[i]}) f(x^{[i-1]}x_i z^{(i,n)}) \right. \\ &\quad \left. - \int_{\mathbf{X}^{(i,n)}} \mu^{(i,n)}(dz^{(i,n)}|x^{[i-1]}y_i) f(x^{[i-1]}y_i z^{(i,n)}) \right) \\ &= \int_{\mathbf{X}^{[i,n]}} \mu^{[i,n]}(dy^{[i,n]}|x^{[i-1]}) \left(K^{(i+1)}f(x^{[i-1]}x_i y^{(i,n)}) - K^{(i+1)}f(x^{[i-1]}y_i y^{(i,n)}) \right). \end{aligned}$$

From this, it follows that $A^{(i)} \leq M^{(i)} \leq B^{(i)}$ a.s., where

$$\begin{aligned} A^{(i)} &\triangleq \int_{\mathbf{X}^{[i,n]}} \mu^{[i,n]}(\mathrm{d}y^{[i,n]} | x^{[i-1]}) \inf_{x_i \in \mathbf{X}_i} \left(K^{(i+1)} f(x^{[i-1]} x_i y^{(i,n)}) - K^{(i+1)} f(x^{[i-1]} y_i y^{(i,n)}) \right) \\ B^{(i)} &\triangleq \int_{\mathbf{X}^{[i,n]}} \mu^{[i,n]}(\mathrm{d}y^{[i,n]} | x^{[i-1]}) \sup_{x_i \in \mathbf{X}_i} \left(K^{(i+1)} f(x^{[i-1]} x_i y^{(i,n)}) - K^{(i+1)} f(x^{[i-1]} y_i y^{(i,n)}) \right), \end{aligned}$$

and

$$\|B^{(i)} - A^{(i)}\|_\infty \leq \|\rho_i\| \delta_i(K^{(i+1)} f). \quad (4.22)$$

By definition of the Wasserstein matrix, we have

$$\delta_i(K^{(i+1)} f) \leq \sum_{j \in T} V_{ij}^{(i+1)} \delta_j(f).$$

Substituting this estimate into (4.22), we get

$$\sum_{i=1}^n \|B^{(i)} - A^{(i)}\|_\infty^2 \leq \sum_{i=1}^n |(\Gamma \delta(f))_i|^2 \equiv \|\Gamma \delta(f)\|_{\ell^2(T)}^2. \quad (4.23)$$

The probability estimate (4.21) then follows from the Azuma–Hoeffding inequality (4.19). \blacksquare

We can also use the martingale method to obtain a tensorization result for transportation inequalities without independence assumptions. This result, which generalizes a theorem of Djellout, Guillin, and Wu [8, Thm. 2.11], can be used even when the metrics ρ_i are not necessarily bounded.

Theorem 4.3 *Suppose that there exist constants $c_1, \dots, c_n \geq 0$, such that*

$$\mu_i(\cdot | x^{[i-1]}) \in T_{\rho_i}(c_i), \quad \forall i \in T, x^{[i-1]} \in \mathbf{X}^{[i-1]}. \quad (4.24)$$

For each $i \in \{1, \dots, n+1\}$, let $V^{(i)}$ be a Wasserstein matrix for $K^{(i)}$. Then $\mu \in T_\rho(c)$ with

$$c = \sum_{i \in T} c_i \left(\sum_{j \in T} V_{ij}^{(i+1)} \right)^2. \quad (4.25)$$

Proof: By the Bobkov–Götze theorem [1], it suffices to show that, for every $f : \mathbf{X} \rightarrow \mathbb{R}$ with $\|f\|_{\text{Lip}} \leq 1$, the random variable $f(X)$ with $X \sim \mu$ is c -subgaussian, with c given by (4.25). To that end, we again consider the martingale decomposition

$$f - \mathbb{E}_\mu[f] = \sum_{i \in T} M^{(i)}$$

with $M^{(i)} = K^{(i+1)} f - K^{(i)} f$. We will show that, for every i ,

$$\log \mathbb{E}_\mu \left[e^{\lambda M^{(i)}} \middle| X^{[i-1]} \right] \leq \frac{c_i \left(\sum_{j \in T} V_{ij}^{(i+1)} \right)^2 \lambda^2}{2}, \quad \forall \lambda \in \mathbb{R}. \quad (4.26)$$

This, in turn, will yield the desired subgaussian estimate

$$\begin{aligned}\mathbb{E}_\mu \left[e^{\lambda(f - \mathbb{E}_\mu[f])} \right] &= \mathbb{E}_\mu \left[\exp \left(\lambda \sum_{i \in T} M^{(i)} \right) \right] \\ &\leq \exp \left(\frac{c\lambda^2}{2} \right)\end{aligned}$$

for every $\lambda \in \mathbb{R}$.

To proceed, note that, for a fixed realization $x^{[i-1]}$ of $X^{[i-1]}$, $M^{(i)} = K^{(i+1)}f - K^{(i)}f$ is $\sigma(X_i)$ -measurable, and

$$\begin{aligned}\|M^{(i)}\|_{\text{Lip}} &\leq \sup_{\substack{x, y \in \mathsf{X} \\ x^{T \setminus \{i\}} = y^{T \setminus \{i\}}}} \frac{|K^{(i+1)}f(x) - K^{(i+1)}f(y)|}{\rho_i(x_i, y_i)} \\ &\equiv \delta_i(K^{(i+1)}f) \\ &\leq \sum_{j \in T} V_{ij}^{(i+1)} \delta_j(f) \\ &\leq \sum_{j \in T} V_{ij}^{(i+1)},\end{aligned}$$

where we have used the definition of the Wasserstein matrix, as well as the fact that $\|f\|_{\text{Lip}} \leq 1$ is equivalent to $\delta_j(f) \leq 1$ for all $j \in T$. Since $\mu_i(\cdot | x^{[i-1]}) \in T_{\rho_i}(c)$ by hypothesis, we obtain the estimate (4.26) by the Bobkov–Götze theorem. \blacksquare

As a sanity check, let us confirm that, in the case when μ is a product measure and the product space X is endowed with the weighted Hamming metric ρ_α defined in (2.3), Theorems 4.2 and 4.3 both reduce to McDiarmid’s inequality. To see this, we first note that, when the X_i ’s are independent, we can write

$$K^{(i)}f(x) = \int_{\mathsf{X}^{[i, n]}} f(x^{[i-1]}y^{[i, n]}) \mu_i(dy_i) \mu_{i+1}(dy_{i+1}) \dots \mu_n(dy_n)$$

for each $i \in T$, $f \in L^1(\mu)$, and $x \in \mathsf{X}$. This, in turn, implies that

$$\begin{aligned}\delta_i(K^{(i+1)}f) &= \alpha_i^{-1} \sup_{\substack{x, z \in \mathsf{X} \\ x^{T \setminus \{i\}} = z^{T \setminus \{i\}}}} \left| K^{(i+1)}f(x) - K^{(i+1)}f(z) \right| \\ &= \alpha_i^{-1} \sup_{\substack{x, z \in \mathsf{X} \\ x^{T \setminus \{i\}} = z^{T \setminus \{i\}}}} \left| \int_{\mathsf{X}^{(i, n]}} f(x^{[i]}y^{(i, n]}) \mu_{i+1}(dy_{i+1}) \dots \mu_n(dy_n) \right. \\ &\quad \left. - \int_{\mathsf{X}^{(i, n]}} f(z^{[i]}y^{(i, n]}) \mu_{i+1}(dy_{i+1}) \dots \mu_n(dy_n) \right| \\ &\leq \alpha_i^{-1} \sup_{\substack{x, z \in \mathsf{X} \\ x^{T \setminus \{i\}} = z^{T \setminus \{i\}}}} |f(x) - f(z)| \\ &= \delta_i(f),\end{aligned}$$

where we have used the fact that, with $\rho_i(x_i, z_i) = \alpha_i \mathbf{1}\{x_i \neq z_i\}$, $\|\rho_i\| = \alpha_i$ for every $i \in T$. Therefore, for each $i \in T$, we can always choose a Wasserstein matrix $V^{(i+1)}$ for $K^{(i+1)}$ in such a

way that its i th row has zeroes everywhere except for the i th column, where it has a 1. Now, for any function $f : \mathsf{X} \rightarrow \mathbb{R}$ which is 1-Lipschitz with respect to ρ_α , we can take $\delta(f) = \mathbf{1}$. Therefore, for any such f Theorem 4.2 gives

$$\mathbb{P}_\mu \left\{ |f(X) - \mathbb{E}_\mu[f(X)]| \geq t \right\} \leq 2 \exp \left(-\frac{2t^2}{\sum_{i=1}^n \alpha_i^2} \right), \quad \forall t \geq 0$$

which is precisely McDiarmid's inequality. Since the constant 2 in McDiarmid's inequality is known to be sharp, this shows that the coefficient 2 in the exponent in (4.21) is likewise optimal. Moreover, with our choice of ρ_i , condition (4.24) of Theorem 4.3 holds with $c_i = \alpha_i^2/4$, and, in light of the discussion above, we can arrange $\sum_j V_{ij}^{(i+1)} = 1$. Therefore, by Theorem 4.3, any function $f : \mathsf{X} \rightarrow \mathbb{R}$ which is 1-Lipschitz with respect to ρ_α is c -subgaussian with constant

$$c = \sum_{i \in T} c_i \left(\sum_{j \in T} V_{ij}^{(i+1)} \right)^2 = \frac{1}{4} \sum_{i \in T} \alpha_i^2,$$

which is just another equivalent statement of McDiarmid's inequality.

It is also possible to consider alternative choices of the filtration $\{\mathcal{F}_j\}_{j=0}^m$. For example, if we partition the index set T into m disjoint subsets (blocks) T_1, \dots, T_m , we can take

$$\mathcal{F}_j \triangleq \sigma(X_i : i \in \Lambda_j), \quad \forall i \in T$$

where $\Lambda_j \triangleq T_1 \cup \dots \cup T_j$. Defining for each $j \in [m]$ the Markov kernel

$$\tilde{K}^{(j)}(x, dy) \triangleq \delta_{x^{\Lambda_{i-1}}} \otimes \mu^{T \setminus \Lambda_{i-1}}(dy^{T \setminus \Lambda_{i-1}} | x^{\Lambda_{i-1}}),$$

we can write

$$M^{(j)} = \mathbb{E}_\mu[f(X) | \mathcal{F}_j] - \mathbb{E}_\mu[f(X) | \mathcal{F}_{j-1}] = K^{(j+1)}f - K^{(j)}f$$

for every $j \in [m]$. As before, we take $K^{(1)}f = \mathbb{E}_\mu[f]$ and $K^{(m+1)}f = f$. Given a measurable function $f : \mathsf{X} \rightarrow \mathbb{R}$, we can define the oscillation of f in the j th block T_j , $j \in [m]$, by

$$\tilde{\delta}_j(f) \triangleq \sup_{\substack{x, z \in \mathsf{X} \\ x^{T \setminus T_j} = z^{T \setminus T_j}}} \frac{|f(x) - f(z)|}{\rho^{T_j}(x^{T_j}, z^{T_j})}.$$

The definition of a Wasserstein matrix is modified accordingly: we say that a nonnegative matrix $\tilde{V} = (\tilde{V}_{jk})_{j, k \in [m]}$ is a Wasserstein matrix for a Markov kernel K on X with respect to the partition $\{T_j\}_{j=1}^m$ if, for any Lipschitz function $f : \mathsf{X} \rightarrow \mathbb{R}$,

$$\tilde{\delta}_j(Kf) \leq \sum_{k=1}^m \tilde{V}_{jk} \tilde{\delta}_k(f)$$

for all $j \in [m]$. With these definitions at hand, the following theorem, which generalizes a result of Paulin [29, Thm.2.1], can be proved in the same way as Theorem 4.2:

Theorem 4.4 For each $j \in [m+1]$, let $\tilde{V}^{(j)} = (\tilde{V}_{k\ell}^{(j)})_{k,\ell \in [m]}$ be a Wasserstein matrix for $\tilde{K}^{(j)}$ with respect to the partition $\{T_j\}$. Define the matrix $\tilde{\Gamma} = (\tilde{\Gamma}_{k\ell})_{k,\ell \in [m]}$ with entries

$$\tilde{\Gamma}_{k\ell} \triangleq \|\rho^{T_k}\| \tilde{V}_{k\ell}^{(k+1)},$$

where

$$\|\rho^{T_k}\| \triangleq \sup_{x^{T_k}, z^{T_k}} \rho^{T_k}(x^{T_k}, z^{T_k})$$

is the diameter of the metric space (X^{T_k}, ρ^{T_k}) . Then, for any $f \in \text{Lip}(X, \rho)$ and for any $t \geq 0$,

$$\mathbb{P}_\mu \left\{ |f(X) - \mathbb{E}_\mu[f(X)]| \geq t \right\} \leq 2 \exp \left(- \frac{2t^2}{\|\tilde{\Gamma} \tilde{\delta}(f)\|_{\ell^2(m)}^2} \right).$$

5 The martingale method in action

We now show that several previously published concentration inequalities for functions of dependent random variables arise as special cases of Theorem 4.2 by exploiting the freedom to choose the Wasserstein matrices $V^{(i)}$. In fact, careful examination of the statement of Theorem 4.2 shows that, for each $i \in T$, we only need to extract the i th row of $V^{(i+1)}$.

5.1 Concentration inequalities under the Dobrushin uniqueness condition

One particularly clean way of constructing the desired Wasserstein matrices is via the classical comparison theorem of Dobrushin for Gibbs measures [9]. For our purposes, we give its formulation due to Föllmer [11]:

Lemma 5.1 Let ν and $\tilde{\nu}$ be two Borel probability measures on X . Define the matrix $C^\nu = (C_{ij}^\nu)_{i,j \in T}$ and the vector $b^{\nu, \tilde{\nu}} = (b_i^{\nu, \tilde{\nu}})_{i \in T}$ by

$$C_{ij}^\nu = \sup_{\substack{x, z \in X \\ x^{T \setminus \{j\}} = z^{T \setminus \{j\}}} \frac{W_i(\nu_i(\cdot | x^{T \setminus \{i\}}), \nu_i(\cdot | z^{T \setminus \{i\}}))}{\rho_j(x_j, z_j)} \quad (5.27)$$

and

$$b_i^{\nu, \tilde{\nu}} = \int_{X^{T \setminus \{i\}}} \tilde{\nu}^{T \setminus \{i\}}(dx^{T \setminus \{i\}}) W_i(\nu_i(\cdot | x^{T \setminus \{i\}}), \tilde{\nu}_i(\cdot | x^{T \setminus \{i\}})). \quad (5.28)$$

Suppose that the spectral radius of C^ν is strictly smaller than unity. Then, for any $f \in L^1(\mu)$,

$$|\mathbb{E}_\nu f - \mathbb{E}_{\tilde{\nu}} f| \leq \sum_{j,k \in T} \delta_j(f) D_{jk}^\nu b_k^{\nu, \tilde{\nu}}, \quad (5.29)$$

where $D^\nu \triangleq \sum_{m=0}^{\infty} (C^\nu)^m$.

Remark 5.1 The matrix C^ν is called the *Dobrushin interdependence matrix* of ν . When the spectral radius of C^ν is strictly smaller than unity, we say that ν satisfies the *Dobrushin uniqueness condition*. This condition is used in statistical physics to establish the absence of phase transitions, which is equivalent to uniqueness of a global Gibbs measure consistent with a given local specification (see the book of Georgii [12] for details).

Given an index $i \in T$, we will extract the i th row of a Wasserstein matrix for $K^{(i+1)}$ by applying the Dobrushin comparison theorem to a particular pair of probability measures on \mathbf{X} . Let $x, z \in \mathbf{X}$ be two configurations that differ only in the i th coordinate: $x^{T \setminus \{i\}} = z^{T \setminus \{i\}}$ and $x_i \neq z_i$. Thus, we can write $z = x^{[i-1]} z_i x^{(i,n)}$, and

$$\begin{aligned}
& K^{(i+1)} f(x) - K^{(i+1)} f(z) \\
&= \mathbb{E}_\mu[f(X) | X^{[i]} = x^{[i-1]} x_i] - \mathbb{E}_\mu[f(X) | X^{[i]} = x^{[i-1]} z_i] \\
&= \int_{\mathbf{X}^{(i,n)}} f(x^{[i-1]} x_i y^{(i,n)}) \mu^{(i,n)}(\mathrm{d}y^{(i,n)} | x^{[i-1]} x_i) - \int_{\mathbf{X}^{(i,n)}} f(x^{[i-1]} z_i y^{(i,n)}) \mu^{(i,n)}(\mathrm{d}y^{(i,n)} | x^{[i-1]} z_i) \\
&= \int_{\mathbf{X}^{(i,n)}} \left(f(x^{[i-1]} x_i y^{(i,n)}) - f(x^{[i-1]} z_i y^{(i,n)}) \right) \mu^{(i,n)}(\mathrm{d}y^{(i,n)} | x^{[i-1]} x_i) \\
&\quad + \int_{\mathbf{X}^{(i,n)}} f(x^{[i-1]} z_i y^{(i,n)}) \mu^{(i,n)}(\mathrm{d}y^{(i,n)} | x^{[i-1]} x_i) - \int_{\mathbf{X}^{(i,n)}} f(x^{[i-1]} z_i y^{(i,n)}) \mu^{(i,n)}(\mathrm{d}y^{(i,n)} | x^{[i-1]} z_i).
\end{aligned} \tag{5.30}$$

By definition of the local oscillation, the first integral in (5.30) is bounded by $\delta_i(f) \rho_i(x_i, z_i)$. To handle the remaining terms, define two probability measures $\nu, \tilde{\nu}$ on \mathbf{X} by

$$\begin{aligned}
\nu(\mathrm{d}y) &\triangleq \delta_{x^{[i-1]} z_i}(\mathrm{d}y^{[i]}) \otimes \mu^{(i,n)}(\mathrm{d}y^{(i,n)} | x^{[i-1]} x_i) \\
\tilde{\nu}(\mathrm{d}y) &\triangleq \delta_{x^{[i-1]} z_i}(\mathrm{d}y^{[i]}) \otimes \mu^{(i,n)}(\mathrm{d}y^{(i,n)} | x^{[i-1]} z_i).
\end{aligned}$$

Using this definition and Lemma 5.1, we can write

$$\begin{aligned}
& \int_{\mathbf{X}^{(i,n)}} f(x^{[i-1]} z_i y^{(i,n)}) \mu^{(i,n)}(\mathrm{d}y^{(i,n)} | x^{[i-1]} x_i) - \int_{\mathbf{X}^{(i,n)}} f(x^{[i-1]} z_i y^{(i,n)}) \mu^{(i,n)}(\mathrm{d}y^{(i,n)} | x^{[i-1]} z_i) \\
&= \int f \mathrm{d}\nu - \int f \mathrm{d}\tilde{\nu} \\
&\leq \sum_{j,k \in T} \delta_j(f) D_{jk}^\nu b_k^{\nu, \tilde{\nu}}.
\end{aligned} \tag{5.31}$$

It remains to obtain explicit upper bounds on the entries of D^ν and $b^{\nu, \tilde{\nu}}$. To that end, we first note that, for a given $j \in T$ and for any $u, y \in \mathbf{X}$,

$$W_j(\nu_j(\cdot | u^{T \setminus \{j\}}), \nu_j(\cdot | y^{T \setminus \{j\}})) = \begin{cases} 0, & j \leq i \\ W_j(\mu_j(\cdot | x^{[i-1]} x_i u^{(i,n) \setminus \{j\}}), \mu_j(\cdot | x^{[i-1]} z_i u^{(i,n) \setminus \{j\}})), & j > i \end{cases}$$

Therefore, $C_{jk}^\nu \leq C_{jk}^\mu$. Likewise, for a given $k \in T$ and for any $y \in \mathbf{X}$,

$$W_k(\nu_k(\cdot | y^{T \setminus \{k\}}), \tilde{\nu}_k(\cdot | y^{T \setminus \{k\}})) = \begin{cases} 0, & k \leq i \\ W_k(\mu_k(\cdot | x^{[i-1]} x_i y^{(i,n) \setminus \{k\}}), \mu_k(\cdot | x^{[i-1]} z_i y^{(i,n) \setminus \{k\}})), & k > i \end{cases}$$

Therefore, $b_k^{\nu, \bar{\nu}} \leq C_{ki}^\mu \rho_i(x_i, z_i)$. Since the matrices C^ν and C^μ are nonnegative, $D_{jk}^\nu \leq D_{jk}^\mu$. Consequently, we can write

$$\begin{aligned} \int f d\nu - \int f d\bar{\nu} &\leq \sum_{j,k \in T} \delta_j(f) D_{jk}^\mu C_{ki}^\mu \rho_i(x_i, z_i) \\ &= \sum_{j \in T} \delta_j(f) (D^\mu C^\mu)_{ji} \rho_i(x_i, z_i) \\ &= \sum_{j \in T} \delta_j(f) (D^\mu - \text{id})_{ji} \rho_i(x_i, z_i). \end{aligned} \quad (5.32)$$

Therefore, from (5.31) and (5.32), we have

$$\begin{aligned} \frac{K^{(i+1)} f(x) - K^{(i+1)} f(z)}{\rho_i(x_i, z_i)} &\leq \delta_i(f) + \sum_{j \in T} (D^\mu - \mathbf{1})_{ij}^\top \delta_j(f) \\ &= \sum_{j \in T} (D^\mu)_{ij}^\top \delta_j(f). \end{aligned} \quad (5.33)$$

We have thus proved the following:

Corollary 5.1 *Suppose that the probability measure μ satisfies the Dobrushin uniqueness condition, i.e., the spectral radius of its Dobrushin interdependence matrix C^μ is strictly smaller than unity. Then, for any $t \geq 0$, the concentration inequality (4.21) holds with*

$$\Gamma_{ij} = \|\rho_i\| (D^\mu)_{ij}^\top, \quad i, j \in T. \quad (5.34)$$

For example, when each X_i is equipped with the trivial metric $\rho_i(x_i, z_i) = \mathbf{1}\{x_i \neq z_i\}$, we have $\|\rho_i\| = 1$ for all i , and consequently obtain the concentration inequality

$$\mathbb{P}_\mu \left\{ |f(X) - \mathbb{E}_\mu[f(X)]| \geq t \right\} \leq 2 \exp \left(- \frac{2t^2}{\|(D^\mu)^\top \boldsymbol{\delta}(f)\|_{\ell^2(T)}^2} \right). \quad (5.35)$$

The same inequality, but with a worse constant in the exponent, was obtained by Külske [19, p. 45].

5.2 Concentration inequalities via couplings

Another method for constructing Wasserstein matrices for the Markov kernels $K^{(i)}$ is via couplings. One notable advantage of this method is that it does not explicitly rely on the Dobrushin uniqueness condition; however, some such condition is typically necessary in order to obtain good bounds for the norm $\|\Gamma \boldsymbol{\delta}(f)\|_{\ell^2(T)}$.

Fix an index $i \in T$ and any two $x, z \in \mathsf{X}$ that differ only in the i th coordinate: $x^{T \setminus \{i\}} = z^{T \setminus \{i\}}$ and $x_i \neq z_i$. Let $\mathbf{P}_{x,z}^{[i]}$ be any coupling of the conditional laws $\mu^{(i,n]}(\cdot | x^{[i]})$ and $\mu^{(i,n]}(\cdot | z^{[i]})$. Then for any $f \in L^1(\mu)$ we can write

$$\begin{aligned} &K^{(i+1)} f(x) - K^{(i+1)} f(z) \\ &= \int_{\mathsf{X}^{(i,n]} \times \mathsf{X}^{(i,n]}} \mathbf{P}_{x,z}^{[i]}(du^{(i,n]}, dy^{(i,n]}) \left(f(x^{[i]}, u^{(i,n]}) - f(z^{[i]}, y^{(i,n]}) \right) \\ &\leq \delta_i(f) \rho_i(x_i, z_i) + \sum_{j \in T: j > i} \delta_j(f) \int_{\mathsf{X}^{(i,n]} \times \mathsf{X}^{(i,n]}} \mathbf{P}_{x,z}^{[i]}(du^{(i,n]}, dy^{(i,n]}) \rho_j(u_j, y_j). \end{aligned}$$

Therefore,

$$\begin{aligned} \frac{|K^{(i+1)}f(x) - K^{(i+1)}f(z)|}{\rho_i(x_i, z_i)} &\leq \delta_i(f) + \sum_{j \in T: j > i} \frac{\int \rho_j d\mathbf{P}_{x,z}^{[i]}}{\rho_i(x_i, z_i)} \delta_j(f) \\ &\leq \delta_i(f) + \sum_{j \in T: j > i} \sup_{\substack{x, z \in \mathbf{X} \\ x^{T \setminus \{i\}} = z^{T \setminus \{i\}}} \frac{\int \rho_j d\mathbf{P}_{x,z}^{[i]}}{\rho_i(x_i, z_i)} \delta_j(f). \end{aligned}$$

Remembering that we only need the i th row of a Wasserstein matrix for $K^{(i+1)}$, we may take

$$V_{ij}^{(i+1)} = \begin{cases} 0, & j < i \\ 1, & j = i \\ \sup_{\substack{x, z \in \mathbf{X} \\ x^{T \setminus \{i\}} = z^{T \setminus \{i\}}} \frac{\int \rho_j d\mathbf{P}_{x,z}^{[i]}}{\rho_i(x_i, z_i)}, & j > i \end{cases}. \quad (5.36)$$

We have thus proved the following:

Corollary 5.2 *For each index $i \in T$ and for each pair $x, z \in \mathbf{X}$ of configurations with $x^{T \setminus \{i\}} = z^{T \setminus \{i\}}$, pick an arbitrary coupling $\mathbf{P}_{x,z}^{[i]}$ of the conditional laws $\mu^{\{i\}}(\cdot | x^{[i]})$ and $\mu^{\{i\}}(\cdot | z^{[i]})$. Then, for any $t \geq 0$, the concentration inequality (4.21) holds with*

$$\Gamma_{ij} = \|\rho_i\| V_{ij}^{(i+1)}, \quad i, j \in T \quad (5.37)$$

where the entries $V_{ij}^{(i+1)}$ are given by (5.36).

In the case when each \mathbf{X}_i is equipped with the trivial metric $\rho_i(x_i, z_i) = \mathbf{1}\{x_i \neq z_i\}$, the entries Γ_{ij} for $j > i$ take the form

$$\Gamma_{ij} = \sup_{\substack{x, z \in \mathbf{X} \\ x^{T \setminus \{i\}} = z^{T \setminus \{i\}}} \mathbf{P}_{x,z}^{[i]} \{Y_j^{(0)} \neq Y_j^{(1)}\}, \quad (5.38)$$

where $(Y^{(0)}, Y^{(1)}) = ((Y_{i+1}^{(0)}, \dots, Y_n^{(0)}), (Y_{i+1}^{(1)}, \dots, Y_n^{(1)}))$ is a random object taking values in $\mathbf{X}^{(i,n]} \times \mathbf{X}^{(i,n]}$. A special case of this construction, under the name of *coupling matrix*, was used by Chazottes et al. [6]. In that work, each $\mathbf{P}_{x,z}^{[i]}$ was chosen to minimize

$$\mathbf{P}\{Y^{(0)} \neq Y^{(1)}\},$$

over all couplings \mathbf{P} of $\mu^{(i,n]}(\cdot | x^{[i]})$ and $\mu^{(i,n]}(\cdot | z^{[i]})$, in which case we have

$$\begin{aligned} \mathbf{P}_{x,z}^{[i]} \{Y^{(0)} \neq Y^{(1)}\} &= \inf_{\mathbf{P} \in \mathcal{C}(\mu^{(i,n]}(\cdot | x^{[i]}), \mu^{(i,n]}(\cdot | z^{[i]}))} \mathbf{P}\{Y^{(0)} \neq Y^{(1)}\} \\ &= \left\| \mu^{(i,n]}(\cdot | x^{[i]}) - \mu^{(i,n]}(\cdot | z^{[i]}) \right\|_{\text{TV}}. \end{aligned}$$

However, it is not clear how to relate the quantities $\mathbf{P}_{x,z}^{[i]} \{Y_j^{(0)} \neq Y_j^{(1)}\}$ and $\mathbf{P}_{x,z}^{[i]} \{Y^{(0)} \neq Y^{(1)}\}$, apart from the obvious bound

$$\mathbf{P}_{x,z}^{[i]} \{Y_j^{(0)} \neq Y_j^{(1)}\} \leq \mathbf{P}_{x,z}^{[i]} \{Y^{(0)} \neq Y^{(1)}\} = \left\| \mu^{(i,n]}(\cdot | x^{[i]}) - \mu^{(i,n]}(\cdot | z^{[i]}) \right\|_{\text{TV}},$$

which gives

$$\Gamma_{ij} \leq \sup_{\substack{x, z \in \mathbf{X} \\ x^{T \setminus \{i\}} = z^{T \setminus \{i\}}}} \left\| \mu^{(i, n]}(\cdot | x^{[i]}) - \mu^{(i, n]}(\cdot | z^{[i]}) \right\|_{\text{TV}}.$$

An alternative choice of coupling is the so-called *maximal coupling* due to Goldstein [13], which for our purposes can be described as follows: let $U = (U_\ell)_{\ell=1}^m$ and $Y = (Y_\ell)_{\ell=1}^m$ be two random m -tuples taking values in a product space $\mathbf{E} = \mathbf{E}_1 \times \dots \times \mathbf{E}_m$, where each \mathbf{E}_ℓ is Polish. Then there exists a coupling \mathbf{P} of the probability laws $\mathcal{L}(U)$ and $\mathcal{L}(Y)$, such that

$$\mathbf{P} \left\{ U^{[\ell, m]} \neq Y^{[\ell, m]} \right\} = \left\| \mathcal{L}(U^{[\ell, m]}) - \mathcal{L}(Y^{[\ell, m]}) \right\|_{\text{TV}}, \quad \ell \in \{1, \dots, m\}. \quad (5.39)$$

Thus, for each $i \in T$ and for every pair $x, z \in \mathbf{X}$ with $x^{T \setminus \{i\}} = z^{T \setminus \{i\}}$, let $\mathbf{P}_{x, z}^{[i]}$ be the Goldstein coupling of $\mu^{(i, n]}(\cdot | x^{[i]})$ and $\mu^{(i, n]}(\cdot | z^{[i]})$. Then for each $j \in \{i+1, \dots, n\}$, using (5.39) we have

$$\begin{aligned} \mathbf{P}_{x, z}^{[i]} \{ Y_j^{(0)} \neq Y_j^{(1)} \} &\leq \mathbf{P}_{x, z}^{[i]} \{ (Y_j^{(0)}, \dots, Y_n^{(0)}) \neq (Y_j^{(1)}, \dots, Y_n^{(1)}) \} \\ &= \left\| \mu^{[j, n]}(\cdot | x^{[i]}) - \mu^{[j, n]}(\cdot | z^{[i]}) \right\|_{\text{TV}}. \end{aligned}$$

This choice of coupling gives rise to the upper-triangular matrix $\Gamma = (\Gamma_{ij})_{i, j \in T}$ with

$$\Gamma_{ij} = \begin{cases} 0, & i < j \\ 1, & i = j \\ \sup_{\substack{x, z \in \mathbf{X} \\ x^{T \setminus \{i\}} = z^{T \setminus \{i\}}}} \left\| \mu^{[j, n]}(\cdot | x^{[i]}) - \mu^{[j, n]}(\cdot | z^{[i]}) \right\|_{\text{TV}}, & i > j \end{cases}. \quad (5.40)$$

Substituting this matrix into (4.21), we recover the concentration inequality of Kontorovich and Ramanan [17], but with an improved constant in the exponent.

Remark 5.2 It was erroneously claimed in [16, 14, 15] that the basic concentration inequalities of Chazottes et al. [6] and Kontorovich and Ramanan [17] are essentially the same, only derived using different methods. As the discussion above elucidates, the two methods use different couplings (the former, explicitly, and the latter, implicitly) — which yield quantitatively different and, in general, incomparable mixing coefficients.

Remark 5.3 Kontorovich and Ramanan obtained the matrix (5.40) using analytic methods without constructing an explicit coupling. In 2012, S. Shlosman posed the following question: could this matrix have been derived using a suitable coupling? We can now answer his question in the affirmative: the coupling is precisely Goldstein's maximal coupling.

As an illustration, let us consider two specific types of the probability law μ : a directed Markov model (i.e., a Markov chain) and an undirected Markov model (i.e., a Gibbsian Markov random field). In the directed case, suppose that the elements of T are ordered in such a way that μ can be disintegrated in the form

$$\mu(dx) = \mu_1(dx_1) \otimes K_1(x_1, dx_2) \otimes K_2(x_2, dx_3) \otimes \dots \otimes K_{n-1}(x_{n-1}, dx_n), \quad (5.41)$$

where μ_0 is a Borel probability measure on $(\mathbf{X}_1, \mathcal{B}_1)$, and, for each $i \in [1, n-1]$, K_i is a Markov kernel from \mathbf{X}_i to \mathbf{X}_{i+1} . For each $i \in [1, n)$, let

$$\theta_i \triangleq \sup_{x_i, z_i \in \mathbf{X}_i} \|K_i(x_i, \cdot) - K_i(z_i, \cdot)\|_{\text{TV}}$$

be the Dobrushin contraction coefficient of K_i . Fix $1 \leq i < j \leq n$ and $x^{[i-1]} \in \mathbf{X}^{[i-1]}$, $y_i, y'_i \in \mathbf{X}_i$. An easy calculation [14] shows that, defining the signed measures η_i on \mathbf{X}_{i+1} by $\eta(\mathrm{d}x_{i+1}) = K_i(y_i, \mathrm{d}x_{i+1}) - K_i(y'_i, \mathrm{d}x_{i+1})$ and ζ_j on \mathbf{X}_j by

$$\zeta_j = \eta_i K_i K_{i+1} K_{i+2} \dots K_{j-1},$$

we have

$$\left\| \mu^{[j,n]}(\cdot | x^{[i-1]} y_i) - \mu^{[j,n]}(\cdot | x^{[i-1]} y'_i) \right\|_{\text{TV}} = \|\zeta_j\|_{\text{TV}} \leq \theta_i \theta_{i+1} \dots \theta_{j-1}, \quad (5.42)$$

where (2.4) was repeatedly invoked to obtain the last inequality. The above yields an upper bound on the Γ_{ij} in (5.40) and hence in the corresponding concentration inequality in (4.21). When more delicate (e.g., spectral [15, 29]) estimates on $\|\zeta_j\|_{\text{TV}}$ are available, these translate directly into tighter concentration bounds.

In the undirected case, μ is a Gibbsian Markov random field induced by pair potentials [12]. To keep things simple, we assume that the local spaces \mathbf{X}_i are all finite. Define an undirected graph with vertex set $T = [n]$ and edge set $E = \{[i, i+1] : 1 \leq i < n\}$ (i.e., a chain graph with vertex set T). Associate with each edge $(i, j) \in E$ a potential function $\psi_{ij} : \mathbf{X}_i \times \mathbf{X}_j \rightarrow [0, \infty)$. Together, these define a probability measure μ on \mathbf{X} via

$$\mu(x) = \frac{\prod_{(i,j) \in E} \psi_{ij}(x_i, x_j)}{\sum_{y \in \mathbf{X}} \prod_{(i,j) \in E} \psi_{ij}(y_i, y_j)}.$$

Since μ is a Markov measure on \mathbf{X} , there is a sequence of Markov kernels K_1, \dots, K_{n-1} generating μ in the sense of (5.41). It is shown in [14] that the contraction coefficient θ_i of the kernel K_i is bounded by

$$\theta_i \leq \frac{R_i - r_i}{R_i + r_i},$$

where

$$R_i = \sup_{(x_i, x_{i+1}) \in \mathbf{X}_i \times \mathbf{X}_{i+1}} \psi_{i,i+1}(x_i, x_{i+1}), \quad r_i = \inf_{(x_i, x_{i+1}) \in \mathbf{X}_i \times \mathbf{X}_{i+1}} \psi_{i,i+1}(x_i, x_{i+1}).$$

The estimate above implies a concentration result, either via (5.42) or via (5.35). To apply the latter, recall that $D^\mu = \sum_{k=1}^{\infty} (C^\mu)^k$, where C^μ is the Dobrushin interdependence matrix defined in (5.27). Assuming that ρ is the unweighted Hamming metric (i.e., $\rho_i(x_i, z_i) = \mathbf{1}\{x_i \neq z_i\}$ for all i) and that the θ_i 's are all majorized by some $\theta < 1$, it is easy to see that $(C^\mu)_{ij} \leq \theta^{|i-j|}$.

6 Open questions

Our focus in this chapter has been on the martingale method for establishing concentration inequalities. In the case of product measures, other techniques, such as the entropy method or

transportation-information inequalities, often lead to sharper bounds. However, these alternative techniques are less developed in the dependent setting, and there appears to be a gap between what is achievable using the martingale method and what is achievable using other means. We close the chapter by listing some open questions that are aimed at closing this gap:

- **(Approximate) tensorization of entropy.** In the independent case, it is possible to derive the same concentration inequality (e.g., McDiarmid’s inequality) using either the martingale method or the entropy method, often with the same sharp constants. However, once the independence assumption is dropped, the situation is no longer so simple. Consider, for example, tensorization of entropy. Several authors (see, e.g., [26, 3, 27]) have obtained so-called *approximate* tensorization inequalities for the relative entropy in the case of weakly dependent random variables: under certain regularity conditions on μ , there exists a constant $A_\mu \geq 1$, such that, for any other probability measure ν ,

$$D(\nu\|\mu) \leq A_\mu \cdot \sum_{i \in T} \mathbb{E}_\nu D(\nu_i(\cdot|X^{T \setminus \{i\}})\|\mu_i(\cdot|X^{T \setminus \{i\}})). \quad (6.43)$$

Having such an inequality in hand, one can proceed to prove concentration for Lipschitz functions in exactly the same way as in the independent case. However, it seems that the constants A_μ in (6.43) are not sharp in the sense that the resulting concentration inequalities are typically worse than what one can obtain using Theorems 4.2 or 4.3 under the same assumptions on μ and f . This motivates the following avenue for further investigation: Derive sharp inequalities of the form (6.43) by relating the constant A_μ to appropriately chosen Wasserstein matrices.

- **General Wasserstein-type matrices.** Using the techniques pioneered by Marton, Samson proved the following concentration of measure result: Consider a function $f : \mathsf{X} \rightarrow \mathbb{R}$ satisfying an “asymmetric” Lipschitz condition of the form

$$f(x) - f(y) \leq \sum_{i \in T} \alpha_i(x) \mathbf{1}\{x_i \neq y_i\}, \quad \forall x, y \in \mathsf{X}$$

for some functions $\alpha_i : \mathsf{X} \rightarrow \mathbb{R}$, such that $\sum_{i \in T} \alpha_i^2(x) \leq 1$ for all $x \in \mathsf{X}$. Then, for any Borel probability measure μ on X , we have

$$\mathbb{P}_\mu \left\{ f(X) - \mathbb{E}_\mu[f(X)] \geq t \right\} \leq \exp \left(-\frac{t^2}{2\|\Delta\|_2^2} \right), \quad (6.44)$$

where the matrix Δ has entries $\Delta_{ij} = \sqrt{\Gamma_{ij}}$ with Γ_{ij} given by (5.40), and

$$\|\Delta\|_2 \triangleq \sup_{v \in \mathbb{R}^T \setminus \{0\}} \frac{\|\Delta v\|_{\ell^2(T)}}{\|v\|_{\ell^2(T)}}$$

is the operator norm of Δ . A more general result in this vein was derived by Marton [24], who showed that an inequality of the form (6.44) holds with Δ computed in terms of any matrix Γ of the form (5.36), where each ρ_i is the trivial metric. Samson’s proof relies on a fairly intricate recursive coupling argument. It would be interesting to develop analogs of (6.44) for arbitrary choices of the metrics ρ_i and with full freedom to choose the Wasserstein matrices $V^{(i)}$ for each $i \in T$. A recent paper by Wintenberger [35] pursues this line of work.

- **The method of exchangeable pairs and Wasserstein matrices.** An alternative route towards concentration inequalities in the dependent setting is via Stein’s method of exchangeable pairs [4, 5]. Using this method, Chatterjee obtained the following result [4, Chap. 4]: Let $f : \mathsf{X} \rightarrow \mathbb{R}$ be a function which is 1-Lipschitz with respect to the weighted Hamming metric ρ_α defined in (2.3). Let μ be a Borel probability measure on X , whose Dobrushin interdependence matrix C^μ satisfies the condition $\|C^\mu\|_2 < 1$. Then, for any $t \geq 0$,

$$\mathbb{P}_\mu \left\{ |f(X) - \mathbb{E}_\mu[f(X)]| \geq t \right\} \leq 2 \exp \left(- \frac{(1 - \|C^\mu\|_2)t^2}{\sum_{i \in T} \alpha_i^2} \right). \quad (6.45)$$

The key ingredient in the proof of (6.45) is the so-called Gibbs sampler, i.e., the Markov kernel \bar{K} on X given by

$$\bar{K}(x, dy) \triangleq \frac{1}{|T|} \sum_{i \in T} \delta_{x^{T \setminus \{i\}}} (dy^{T \setminus \{i\}}) \otimes \mu_i(dy_i | x^{T \setminus \{i\}}).$$

This kernel leaves μ invariant, i.e., $\mu = \mu \bar{K}$, and it is easy to show (see, e.g., [27]) that it contracts the \bar{W} distance: for any other probability measure ν on X ,

$$\bar{W}(\mu \bar{K}, \nu \bar{K}) \leq \left(1 - \frac{1 - \|C^\mu\|_2}{|T|} \right) \bar{W}(\mu, \nu).$$

Since one can obtain contraction estimates for Markov kernels using Wasserstein matrices, it is natural to ask whether Chatterjee’s result can be derived as a special case of a more general method, which would let us freely choose an arbitrary Markov kernel K that leaves μ invariant and control the constants in the resulting concentration inequality by means of a judicious choice of a Wasserstein matrix for K . Such a method would most likely rely on general comparison theorems for Gibbs measures [31].

Acknowledgments

The second author would like to thank IMA for an invitation to speak at the workshop on Information Theory and Concentration Phenomena in Spring 2015, which was part of the annual program “Discrete Structures: Analysis and Applications.” The authors are grateful to the anonymous referee for several constructive suggestions, and to Dr. Naci Saldi for spotting an error in an earlier version of the manuscript. A. Kontorovich was partially supported by the Israel Science Foundation (grant No. 1141/12) and a Yahoo Faculty award. M. Raginsky would like to acknowledge the support of the U.S. National Science Foundation via CAREER award CCF-1254041.

References

- [1] Sergey G. Bobkov and Friedrich Götze. Exponential integrability and transportation cost related to logarithmic Sobolev inequalities. *J. Funct. Anal.*, 163:1–28, 1999.
- [2] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.

- [3] Pietro Caputo, Georg Menz, and Prasad Tetali. Approximate tensorization of entropy at high temperature. *Annales de la Faculté des Sciences de Toulouse Sér. 6*, 24(4):691–716, 2015.
- [4] Sourav Chatterjee. *Concentration inequalities with exchangeable pairs*. PhD thesis, Stanford University, 2005.
- [5] Sourav Chatterjee. Stein’s method for concentration inequalities. *Probability Theory and Related Fields*, 138:305–321, 2007.
- [6] Jean-René Chazottes, Pierre Collet, Christof Külske, and Frank Redig. Concentration inequalities for random fields via coupling. *Probability Theory and Related Fields*, 137(1-2):201–225, 2007.
- [7] Imre Csiszár. Information-type measures of difference of probability distributions and indirect observations. *Studia Sci. Math. Hungar.*, 2:299–318, 1967.
- [8] H. Djellout, A. Guillin, and L. Wu. Transportation cost-information inequalities and applications to random dynamical systems and diffusions. *Ann. Probab.*, 32(3B):2702–2732, 2004.
- [9] Roland L. Dobrushin. Prescribing a system of random variables by conditional distributions. *Theory of Probability and Its Applications*, 15(3):458–486, 1970. Translated from Russian.
- [10] Hans Föllmer. Tail structure of Markov chains on infinite product spaces. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete*, 50:273–285, 1979.
- [11] Hans Föllmer. A covariance estimate for Gibbs measures. *Journal of Functional Analysis*, 46:387–395, 1982.
- [12] Hans-Otto Georgii. *Gibbs Measures and Phase Transitions*, volume 9 of *de Gruyter Studies in Mathematics*. Walter de Gruyter & Co., 2nd edition, 2011.
- [13] Sheldon Goldstein. Maximal coupling. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete*, 46(193-204), 1979.
- [14] Aryeh Kontorovich. Obtaining measure concentration from Markov contraction. *Markov Processes and Related Fields*, 4:613–638, 2012.
- [15] Aryeh Kontorovich and Roi Weiss. Uniform Chernoff and Dvoretzky-Kiefer-Wolfowitz-type inequalities for Markov chains and related processes. *Journal of Applied Probability*, 51:1–14, 2014.
- [16] Aryeh (Leonid) Kontorovich. *Measure Concentration of Strongly Mixing Processes with Applications*. PhD thesis, Carnegie Mellon University, 2007.
- [17] Leonid (Aryeh) Kontorovich and Kavita Ramanan. Concentration Inequalities for Dependent Random Variables via the Martingale Method. *Ann. Probab.*, 36(6):2126–2158, 2008.
- [18] Solomon Kullback. A lower bound for discrimination information in terms of variation. *IEEE Trans. Inform. Theory*, 13:126–127, 1967. Correction, volume 16, p. 652, 1970.
- [19] Christof Külske. Concentration inequalities for functions of Gibbs fields with application to diffraction and random Gibbs measures. *Commun. Math. Phys.*, 239:29–51, 2003.

- [20] Paul Lévy. *Problèmes concrets d'analyse fonctionnelle*. Gauthier-Villars, Paris, 1951. 2d ed.
- [21] Andrei A. Markov. Extension of the law of large numbers to dependent quantities. *Izvestiia Fiz.-Matem. Obsch. Kazan Univ.*, 15:135–156, 1906.
- [22] Katalin Marton. Bounding \bar{d} -distance by informational divergence: a method to prove measure concentration. *Ann. Probab.*, 24(2):857–866, 1996.
- [23] Katalin Marton. Measure concentration for a class of random processes. *Probability Theory and Related Fields*, 110(3):427–439, 1998.
- [24] Katalin Marton. Measure concentration and strong mixing. *Studia Scientiarum Mathematicarum Hungarica*, 19(1-2):95–113, 2003.
- [25] Katalin Marton. Measure concentration for Euclidean distance in the case of dependent random variables. *Ann. Probab.*, 32(3):2526–2544, 2004.
- [26] Katalin Marton. An inequality for relative entropy and logarithmic Sobolev inequalities in Euclidean spaces. *Journal of Functional Analysis*, 264(34-61), 2013.
- [27] Katalin Marton. Logarithmic Sobolev inequalities in discrete product spaces: a proof by a transportation cost distance. arXiv.org preprint 1507.02803, July 2015.
- [28] Bernard Maurey. Construction de suites symétriques. *C. R. Acad. Sci. Paris Sér. A-B* 288, (14):A679–A681, 1979.
- [29] Daniel Paulin. Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electronic Journal of Probability*, 20:1–32, 2015.
- [30] Maxim Raginsky and Igal Sason. *Concentration of Measure Inequalities in Information Theory, Communications, and Coding*. Foundations and Trends in Communications and Information Theory. Now Publishers, 2nd edition, 2014.
- [31] Patrick Rebeschini and Ramon van Handel. Comparison theorems for Gibbs measures. *J. Stat. Phys.*, 157:234–281, 2014.
- [32] Emmanuel Rio. Inégalités de Hoeffding pour les fonctions lipschitziennes de suites dépendantes. *C. R. Acad. Sci. Paris Sér. I Math.*, 330(10):905–908, 2000.
- [33] Paul-Marie Samson. Concentration of measure inequalities for Markov chains and Φ -mixing processes. *Ann. Probab.*, 28(1):416–461, 2000.
- [34] Ramon van Handel. Probability in high dimension. ORF 570 Lecture Notes, Princeton University, June 2014.
- [35] Olivier Wintenberger. Weak transport inequalities and applications to exponential and oracle inequalities. *Electronic Journal of Probability*, 20(114):1–27, 2015.
- [36] Liming Wu. Poincaré and transportation inequalities for Gibbs measures under the Dobrushin uniqueness condition. *Ann. Probab.*, 34(5):1960–1989, 2006.